

Obtaining highly excited eigenstates of the localized XX chain via DMRG-X

Trithip Devakul,¹ Vedika Khemani,² Frank Pollmann,^{3,4} David A. Huse,¹ and S. L. Sondhi¹

¹*Department of Physics, Princeton University, Princeton, NJ 08544, USA*

²*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

³*Max-Planck-Institut für Physik komplexer Systeme, Nöthnitzer Str. 38, 01187 Dresden, Germany*

⁴*Department of Physics, Technical University of Munich, 85748 Garching, Germany*

We benchmark a variant of the recently introduced DMRG-X algorithm against exact results for the localized random field XX chain. We find that the eigenstates obtained via DMRG-X exhibit a highly accurate 1-bit description for system sizes much bigger than the direct, many body, exact diagonalization in the spin variables is able to access. We take advantage of the underlying free fermion description of the XX model to accurately test the strengths and limitations of this algorithm for large system sizes. We discuss the theoretical constraints on the performance of the algorithm from the entanglement properties of the eigenstates, and its actual performance at different values of disorder. Altogether this study offers considerable encouragement for the use of the algorithm even for interacting many body localized spin models that do not possess an underlying free fermion description.

I. INTRODUCTION

The density matrix renormalization group (DMRG) has become one of the most successful numerical tools in the study of ground state properties of one-dimensional systems^{1,2}. DMRG is able to easily go to far larger sizes than accessible by methods such as exact diagonalization, and is even capable of working directly in the thermodynamic limit for translationally invariant systems³. The success of DMRG has been greatly clarified by the advent of quantum information theoretic ideas such as the entanglement structure of quantum states and the language of matrix product states (MPSs)^{2,4}. The basic premise of DMRG is the relatively small entanglement present in the ground states of most gapped local Hamiltonians. Specifically, such states obey an area law, that is, the bipartite entanglement entropy S_E between two subsystems scales with the area of the boundary separating the regions^{5,6}. As a consequence the states can be efficiently represented with controllable error by MPSs which allow efficient computation of desired operator expectation values with controlled errors. This *representability* is coupled in practice with *findability*—which is the existence of efficient algorithms, such as the original DMRG algorithm, which find the desired representation in a time that scales polynomially—indeed linearly—with system size. While for the DMRG algorithm this scaling is an empirical fact about some unknown dynamical system⁷, there now also exists an algorithm that *provably* solves the problem in a time that scales polynomially in system size⁸. DMRG also works very well for critical points, where S_E grows logarithmically with subsystem size which is modest enough to allow sufficient accuracy.

For thermalizing systems this cluster of ideas fails when addressed to highly excited eigenstates—eigenstates with a finite energy density corresponding to a nonzero temperature—due to the volume law entanglement that is generically present in these states. However in many body localized (MBL) systems^{9–14} even highly excited

eigenstates exhibit area law entanglement^{11,15} which leads to the possibility^{16,17} that excited states can be constructed efficiently via DMRG-like algorithms in a manner analogous to ground states. Being able to evade the size restrictions on many body exact diagonalization (MBED) in this fashion would be particularly useful in studying the eigenstate phase transition^{18–20} from a volume-law obeying thermalizing phase at low disorder to an area-law obeying localized phase at strong disorder. While this MBL phase transition has been the subject of intense study numerically^{21–30}, there is still much that is not well understood. Even for the “standard” model of MBL¹¹, the random-field Heisenberg model, estimates of the critical disorder from numerical linked cluster expansion³¹ predict a different value than finite size scaling estimates from exact diagonalization studies^{11,32}. Moreover, a recent study³³ has shown that the entanglement properties of eigenstates jump discontinuously at the MBL-to-thermal transition—even while other properties look continuous—and that the critical point has far less entanglement than previously assumed³⁴, thereby lending encouragement for the potential use of DMRG-like methods all the way to the transition to the thermal phase.

Despite the good news on representability stemming from the area law, we note that the obstacles to constructing highly excited eigenstates of MBL systems are still formidable. One of these obstacles arises from the sheer smallness of gaps in the bulk of the many body spectrum which decrease exponentially with system size. Standard double precision arithmetic only allows a binary representation accurate to 52 bits which translates into an average spectral gap in a system of essentially the same number, say 50, spins. For systems larger than 50 spins we should expect to find exact spectral degeneracies which are artifacts and hence eigenvectors which are misleading superpositions of the underlying eigenstates. Indeed, this will start to happen even before size 50 for selected pairs of levels which are closer than average. The second obstacle is the process of optimization in a random system. Almost all work on ground state DMRG

has been done on undisordered systems and based on general experience with disordered systems there is reason to worry that optimization will not work as well in the more complicated landscapes the latter exhibit. Potential obstacles here are the presence of Griffiths regions and many-body resonances that may cause difficulty in finding an eigenstate.

With this caution injected, we note that the past year has nevertheless seen several studies that have tackled the problem of finding MPS representations for highly excited states. The first of these introduced an approach to finding approximations for *all* eigenstates of a MBL Hamiltonian at one go³⁵ and has been the object of more recent elaboration³⁶. The second, which is the basis of present work, is a variant of DMRG aimed to target highly excited eigenstates, called DMRG-X³⁷. Rather than optimizing an MPS towards the ground state of the Hamiltonian, the DMRG-X algorithm iteratively optimizes the MPS towards an eigenstate of the Hamiltonian based on overlap with the state in the previous step. In a highly localized system, where eigenstates look very close to product states, DMRG-X was shown to converge very well starting from a product state as an initial state and reproduce the results of ED for accessible system sizes. In other works^{38–40}, DMRG based techniques have been used to obtain excited eigenstates via energy (rather than overlap) targetting.

In this paper we return to the DMRG-X algorithm with a view to benchmarking its ability to accurately solve for eigenstates of large systems. We now clarify what we mean by accurately and what we mean by large. One can usefully think of a single DMRG step for a many body system as a many body exact diagonalization (MBED) with soft boundaries so it amplifies the size one can study with a diagonalization subroutine of given power; this is what enables ground states of systems with 100 spins to be studied on a laptop where only 14 (say) could be studied by MBED. Correspondingly we wish to understand the amplification achieved by DMRG-X in the middle of the spectrum over MBED. To gauge this meaningfully and even otherwise, we need to be able to assess the accuracy of the states obtained via DMRG-X. For any system we can quantify the accuracy by computing the variance of the energy in the states that we obtain and comparing it to the level spacing; this criterion was used in Ref. 37 at the limits of machine precision for the variance. However it is not known how to translate from a given accuracy of the variance to the accuracy of various quantities of more physical interest, which limits the utility of this metric.

To get around this limitation, in the present paper we study the the random field XX spin-1/2 chain which has an underlying free fermion representation. The free fermion character is not exploited by DMRG-X, which treats it as it does any interacting spin system. However the ability to exactly solve this model numerically for large system sizes in the fermionic version allows us to check the results of DMRG-X to larger system sizes

than are accessible via MBED. Specifically, we make use of a defining characteristic of localized phases, namely the existence of an emergent set of an extensive number of commuting \mathbb{Z}_2 -valued local integrals of motion^{41,42} (often called “l-bits”), and use the free fermion representation to construct the exact l-bits for very large system sizes. We then compute the expectation values of these l-bits in the DMRG-X obtained states, which allows us to directly assess how well these states approximate the exact eigenstates using a more physical (and economical) measure. A further advantage of studying the XX chain is that we are also able to study the representability of the eigenstates to large system sizes via computation of their entanglement entropies. Overall we find that DMRG-X achieves a roughly three-fold amplification over MBED when we require that all l-bits are accurate to at least 90 percent; most are vastly more accurate. While this particular demonstration is specific to the XX chain, we believe this analysis offers much encouragement that we can trust DMRG-X results even for interacting systems. We note that similar optimism comes from the work of Serbyn et al⁴⁰ who examine the structure of the entanglement spectrum for numerically obtained eigenstates.

In the balance of the paper, we start with a quick recapitulation of the DMRG method in the language of MPSs (Sec. II). Following this, we review the basics of the random XX chain and the accuracy measure (Sec. III), discuss representability where we also address the issue of rare Griffiths effects (Sec. IV), turn to findability where we find it important to modify the DMRG-X to use a hybrid metric for the updates (Sec. V), show that the resulting algorithm is accurate enough to go beyond ED (Sec. VI) and end with some concluding remarks (Sec. VII). Appendices describe some technical details on the locality (App. A) and representation (App. B) of the l-bit operators for the XX chain, and some unexpected wrinkles that crop up in studying the closely related problem of localization in quasiperiodic potentials (App. C).

II. BRIEF REVIEW OF MPS/MPO/DMRG FORMALISM

In this section, we briefly recapitulate the standard DMRG algorithm¹ implemented in the language of matrix-product states (MPSs)². This section closely follows the supplementary material in Ref. 37. A general quantum state $|\Psi\rangle$ for a one-dimensional system of L sites can be written in the following matrix-product state (MPS) form:

$$|\Psi\rangle = \sum_{j_1, \dots, j_L} \sum_{0 < \gamma_n \leq \chi_n} B_{\gamma_1}^{[1]j_1} B_{\gamma_1 \gamma_2}^{[2]j_2} \dots B_{\gamma_{L-1}}^{[L]j_L} |j_1, \dots, j_L\rangle. \quad (1)$$

where $|j_n\rangle$ with $j_n = 1, \dots, d$ is a basis of local states at site n (for a spin 1/2 system, $d = 2$ and $|j_n\rangle = |\uparrow\rangle, |\downarrow\rangle$), and the $B^{[n]}$ are rank three tensors (except on the first

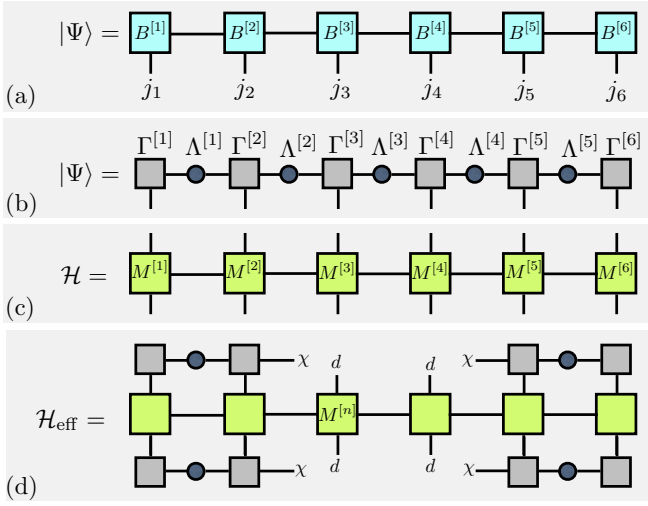


FIG. 1. Diagrammatic representation of (a) the state $|\Psi\rangle$ as an MPS, (b) $|\Psi\rangle$ as a canonical MPS, (c) the Hamiltonian \mathcal{H} as an MPO, and (d) the effective Hamiltonian \mathcal{H}_{eff} in the variational basis.

and last sites where they are rank two tensors). Figure 1(a) shows a pictorial representation of an MPS. The external legs j_n are the “physical” spin indices whereas the internal legs γ_n are the virtual indices that are contracted. Each $B^{[n]j_n}$ is a $\chi_n \times \chi_{n+1}$ matrix (at the boundaries $\chi_1 = \chi_{L+1} = 1$) and each matrix product $\prod_i B^{[i]j_i}$ in Eq. (1) produces a complex number which is the amplitude of $|\Psi\rangle$ on the basis state $|j_1 \cdots j_L\rangle$.

The maximum dimension χ of the $\{B^{[n]}\}$ matrices is called the bond-dimension of the MPS and low entanglement states can be efficiently represented by MPSs of bond dimension $\chi \ll d^{L/2}$. The relationship between χ and the entanglement can be made more precise by considering the Schmidt decomposition of the state $|\Psi\rangle$. For a given bipartition of the system into left and right halves, a singular value decomposition can be used to rewrite $|\Psi\rangle = \sum_{\alpha} \Lambda_{\alpha} |\alpha\rangle_L |\alpha\rangle_R$ where the $|\alpha\rangle_{L/R}$ form orthonormal bases for the left and right halves respectively, and the entanglement entropy of the bipartition is defined through the Schmidt values Λ_{α} as $S_E = -\sum_{\alpha} |\Lambda_{\alpha}|^2 \ln |\Lambda_{\alpha}|^2$. Following a prescription by Vidal⁴³, it is possible to define a *canonical form* (Fig. 1(b)) for the MPS by rewriting each matrix $B^{[n]j_n}$ as a product of a $\chi_n \times \chi_{n+1}$ dimensional complex matrix $\Gamma^{[n]j_n}$ and a square diagonal matrix $\Lambda^{[n]}$ such that matrices $\Lambda^{[n]}$ contain the non-zero Schmidt values for a bipartition between sites n and $n+1$

$$\begin{aligned}
 |\Psi\rangle &= \sum_{j_1, \dots, j_L} \Gamma^{[1]j_1} \Lambda^{[1]} \Gamma^{[2]j_2} \Lambda^{[2]} \dots \Lambda^{[L-1]} \Gamma^{[L]j_L} |j_1, \dots, j_L\rangle \\
 &\equiv \sum_{\alpha=1}^{\chi_{n+1}} \Lambda_{\alpha}^{[n]} |\alpha_n\rangle_L |\alpha_n\rangle_R,
 \end{aligned} \quad (2)$$

and the states $|\alpha_n\rangle_L, |\alpha_n\rangle_R$ define the orthonormal Schmidt states for the left and right halves of the bipar-

tion respectively. This canonical form clearly relates the bond dimension χ to the number of Schmidt values contributing significantly to the entanglement entropy.

Starting from an initial random MPS, the DMRG algorithm iteratively finds the ground state $|\psi_0\rangle$ by sweeping through the system and variationally optimizing the MPS matrices $B^{[n]j_n}$ on neighboring sites to locally minimize the energy $\langle \psi_0 | \mathcal{H} | \psi_0 \rangle$ (keeping the rest of the chain fixed). In the commonly used two-site update which simultaneously updates two sets of matrices $B^{[n]j_n}$ and $B^{[n+1]j_{n+1}}$, an effective Hamiltonian \mathcal{H}_{eff} is constructed by projecting \mathcal{H} to a mixed $\chi_n \chi_{n+2} d^2$ dimensional basis. Here, the local basis states $|j_n\rangle |j_{n+1}\rangle$ represent the two updated sites, and the eigenstates of the reduced density matrix $|\chi_n\rangle_L |\chi_{n+2}\rangle_R$ compactly represent the environment to the left and right of the updated sites. The ground state of \mathcal{H}_{eff} is then found — which is the optimal state for minimizing $\langle \psi_0 | \mathcal{H} | \psi_0 \rangle$ in this subspace — and the matrices on sites $n, n+1$ are updated. The procedure is repeated for all sites until convergence is achieved. The matrix-product operator (MPO) representation of \mathcal{H} , defined exactly analogously to Eq. 1 but now using 4-index tensors M , is shown in Fig. 1(c), and the effective Hamiltonian is depicted pictorially in Fig. 1(d).

The only difference between the ground-state DMRG algorithm outlined above and the DMRG-X algorithm of Ref. 37 is in the update step. In the DMRG-X algorithm, *all* the $d^2 \chi^2$ eigenstates of \mathcal{H}_{eff} are obtained instead of just its ground state, and the matrices $B^{[n]j_n}$ are updated using the eigenstate of \mathcal{H}_{eff} with the maximum overlap with the previously found state in the iterative scheme. As will be discussed in Sec. IV, we will present a modification to this algorithm at this step, where instead of simply choosing the eigenstate with maximal overlap, we also minimize the energy variance within a small subspace of states with high overlap. The algorithm is initialized with an appropriate initial state which is perturbatively “close” to the true eigenstates of the MBL Hamiltonian, such as a product state.

III. THE XX CHAIN

The model we study using DMRG-X is the spin-1/2 XX chain with random fields,

$$\mathcal{H} = - \sum_{i=1}^{L-1} (S_i^x S_{i+1}^x + S_i^y S_{i+1}^y) + \sum_{i=1}^L h_i S_i^z \quad (3)$$

where $S_i^{x,y,z}$ are spin-1/2 operators on site i and $h_i \in [-W, W]$ is chosen from a uniform distribution with width $2W$. This model is Anderson localized for any nonzero W , and the localization length scales as $\xi \approx 25/W^2$ at low W ⁴⁴, while at large W the locator expansion yields $\xi \sim 1/\log W$. This model can be mapped via a Jordan-Wigner transformation onto a system of non-interacting fermions and solved exactly to very large system sizes, which we now proceed to do.

Using the standard Jordan-Wigner substitutions for $S_i^\pm = S_i^x \pm iS_i^y$,

$$S_i^+ = (-1)^{\sum_{j=1}^{i-1} c_j^\dagger c_j} c_i^\dagger, \quad S_i^- = (-1)^{\sum_{j=1}^{i-1} c_j^\dagger c_j} c_i, \quad (4)$$

we arrive at a quadratic Hamiltonian in terms of the fermionic c_i^\dagger, c_i operators

$$\mathcal{H} = -\frac{1}{2} \sum_{i=1}^{L-1} (c_i^\dagger c_{i+1} + \text{h.c.}) - \sum_{i=1}^L h_i c_i^\dagger c_i \quad (5)$$

$$= \sum_{i,j} c_i^\dagger H_{ij} c_j \quad (6)$$

up to a constant. This single particle Hamiltonian H_{ij} is only of dimension L , and can be diagonalized numerically, $H_{ij} = \sum_\alpha U_{i\alpha} \varepsilon_\alpha U_{\alpha j}^\dagger$, giving

$$\mathcal{H} = \sum_{\alpha=1}^L \varepsilon_\alpha a_\alpha^\dagger a_\alpha \quad (7)$$

up to additive constants, where $a_\alpha^\dagger = \sum_i c_i^\dagger U_{i\alpha}$, and similarly $a_\alpha = \sum_i U_{i\alpha}^\dagger c_i$.

In spin-language, we refer to the $\sigma_i^z = 2c_i^\dagger c_i - 1$ as *physical-bit* (p-bit) operators and $\tau_\alpha^z = 2a_\alpha^\dagger a_\alpha - 1$ as *localized-bit* (l-bit) operators, in accordance with standard MBL nomenclature⁴¹. This non-interacting Hamiltonian then takes the simple form $\mathcal{H} = \sum_\alpha h_\alpha \tau_\alpha^z$, where $h_\alpha = \varepsilon_\alpha/2$. The 2^L eigenstates of \mathcal{H} are then obtained by picking each l-bit to be +1 or -1, corresponding to a filled ($n_\alpha = 1$) or empty ($n_\alpha = 0$) fermionic state respectively, and the many-body fermionic eigenstates are simply $|\{n_\alpha\}\rangle = \prod_\alpha a_\alpha^\dagger n_\alpha |0\rangle$.

It is interesting to note that the fermionic raising and lowering operators a_α^\dagger and a_α are inherently nonlocal in spin language due to the $(-1)^{\sum_i c_i^\dagger c_i}$ chain from the Jordan-Wigner transformation. However, in MBL systems one tends to speak of local bosonic raising and lowering operators τ_α^\pm . Bosonic raising and lowering operators can be constructed via a reverse Jordan-Wigner transformation from the $a_\alpha^\dagger, a_\alpha$ operators, but it is not *a priori* clear that these operators are indeed local.⁴⁵ In App. A we show that, indeed, these bosonic raising and lowering operators can be explicitly shown to be local given that the original l-bits are local.

Our primary interest in the l-bit operators is that they can be used to gauge the accuracy of a DMRG state. An exact eigenstate should have $|\langle \tau_\alpha^z \rangle| = 1$ (in the absence of degeneracy). Deviations from this tell us exactly where DMRG has failed to capture the state, and how badly it has done so. The l-bit operators themselves can be efficiently expressed as matrix product operators (MPO) using an internal bond dimension of only 4 (App. B).

The other measure of accuracy is the total energy variance of the state, $\sigma_E^2 = |\langle H^2 \rangle - \langle H \rangle^2|$, which is also efficient to calculate for an MPS. This should be closer to zero the closer a state is to an eigenstate of the Hamiltonian. Expressed in terms of the l-bit operators, σ_E^2

contains expectation values and correlators of these operators, and is thus a related but complementary measure of accuracy.

IV. REPRESENTABILITY: S_E AND GRIFFITHS EFFECTS

We now turn to the question of how efficiently the MB eigenstates of the XX model can be represented, which entails a study of the entanglement properties of these states. We find that rare Griffiths-like effects cause the typical and worst-case entanglement entropy across the bonds of the system to scale differently with L , which in turn affects the maximum bond-dimension needed to represent these states and the scaling of the computational time of the DMRG-X algorithm. While the algorithm always scales polynomially with L , in practice, the power of the polynomial can get quite large as disorder is lowered.

The entanglement spectrum (ES) of a state across a cut determines how well that state can be represented as an MPS with a finite bond dimension χ on that cut. The ES in the MBL phase has been recently shown to decay as a power law⁴⁰. However, the arguments of Ref 40 only apply to interacting MBL systems.⁴⁶ In the XX model, the spectrum of the single particle entanglement Hamiltonian \mathcal{H}_E is found to have a constant density of states on average, similar to the clean system⁴⁷, which will lead to an entanglement spectrum decaying (on average) as

$$\lambda_n \sim \exp \left[-a(\ln n)^2 \right] \quad (8)$$

for some constant a . This is one factor that differentiates our non-interacting model from an interacting model, making it easier to treat with DMRG-X. For now, we ignore the details of the entanglement spectrum and assume that a bond dimension of $\chi \sim e^{S_E}$ is needed to represent a state with entanglement entropy S_E (which is equivalent to assuming a flat ES).

The entanglement entropy S_E can be obtained exactly for a given many-body free-fermion eigenstate $|\psi\rangle = |\{n_\alpha\}\rangle$ using the correlation matrix method⁴⁷. The idea is that the reduced density matrix for a partition A is represented as a thermal density matrix for an “entanglement Hamiltonian” \mathcal{H}_E , $\rho_A = \text{Tr}_{\bar{A}} \rho = e^{-\mathcal{H}_E}$, where \mathcal{H}_E itself is a non-interacting Hamiltonian whose single particle energy eigenvalues can be found from the matrix of two-point correlation functions of fermionic operators evaluated in the state $|\psi\rangle$. The entanglement entropy $S_E = -\text{Tr} \rho_A \ln \rho_A$ can then be calculated knowing the eigenvalues of ρ_A . This knowledge can also be used to obtain the (many-body) eigenvalues of ρ_A , λ_n , which constitute the entanglement spectrum.⁴⁸

Figure 2(left) shows the mean values of the mid-chain entanglement entropy and the maximum entanglement entropy within the whole chain. The *maximum* entanglement entropy S_E^{max} for a given chain length is important as it will determine the maximum bond dimension




FIG. 2. (left) The mean mid-bond and mean max entanglement entropy as a function of system size L . These correspond to the entanglement on the middle bond and the maximum entanglement across all bonds within a sample, respectively, averaged over many samples. While the mean entanglement quickly saturates to an area-law value, the maximum entanglement within a sample continues increasing due to Griffiths regions. (center) A similar plot of the mean mid and max χ as a function of system size, where χ is chosen such that on each cut, the total cutoff error is less than 10^{-5} . Similar to the entanglement entropy, the mean χ saturates to a constant, while the maximum continues increasing with system size. (right) A 2D histogram over all cuts and samples of bond entanglement and necessary bond dimension χ (assuming a naive truncation of the entanglement spectrum) to represent the state with total error² less than 10^{-5} for systems size $L = 320$ at $W = 1.0$, demonstrating the average exponential dependence of χ on S .

χ needed to accurately capture that state as an MPS. While the mid-chain entanglement entropy saturates to an area law, the maximum entanglement entropy continues to grow with L .

The growth of S_E^{\max} with L can be understood as a type of Griffiths effect. While there is no true delocalization in this model, a finite patch of length ℓ appears delocalized if the localization length within that patch exceeds ℓ : $\xi \geq \ell$. Thus, the finite patch appears to have a critical disorder $W_c(\ell) = 5/\sqrt{\ell}$ (this follows from the expression for the localization length $\xi = 25/W^2$). The probability to get such a patch of length ℓ with every onsite field lying within $W_c(\ell)$ is $P(\ell) = (W_c(\ell)/W)^\ell$. Typically, a system of size L will have $N(\ell) = LP(\ell)$ such patches, and equating $N(\ell) \approx 1$ leads to $\ell \sim \log L / \log \log L$ at leading order. This apparently delocalized patch will have volume law entanglement $S_E^{\max} \sim \ell$, and therefore $S_E^{\max} \sim \log L / \log \log L$, which explains the mostly logarithmic scaling with the slight downwards curvature in the S_E^{\max} plot in Figure 2(left). In particular, this analysis shows that the maximum bond dimension needed to represent a particular state which scales as $\log \chi^{\max} \sim S_E^{\max}$ will scale with L , as opposed to approaching a constant as in the case of ground states of gapped systems. In App. C, we discuss how a similar effect may arise even in some quasiperiodic systems, where there are no Griffiths effects.

This analysis also helps us deduce the scaling of the time needed to carry out one sweep of the DMRG-X. If we run the algorithm with variable bond dimension χ on each bond, the locus of the maximum computational time is determined by the tail of the distribution of entanglement entropies, with the shape of the tail determined by the same Griffiths-like effects that were just discussed. If we take $S_E = \kappa \ell$ to be the entanglement of a thermal

region of length ℓ , assuming that the entanglement on each bond is independent, we have that

$$P(S_E) \sim \left(\frac{W_c(\ell(S_E))}{W} \right)^\ell \sim \exp \left[-\frac{S_E}{\kappa} \ln \frac{W}{W_c(S_E/\kappa)} \right] \quad (9)$$

with $W_c(\ell)$ as defined earlier. We therefore see that the tail of $P(S_E)$ will decay slightly faster than exponential, and decay faster for higher W . However, the computational time required to deal with a cut of entanglement S_E scales exponentially with S_E : taking $\chi = e^{S_E}$, we know that the computational cost of diagonalizing the effective Hamiltonian in a DMRG step (a $d^2 \chi^2 \times d^2 \chi^2$ matrix) scales as $(\chi^2)^3 \sim e^{6S_E}$. Thus, the mean computational time will be dominated by bonds with entanglement where $e^{6S_E} P(S_E)$ is maximized. For low W , this implies that the computational time will be dominated by the largest entanglement across any bond in the sample S_E^{\max} . Thus, the growth of S_E^{\max} with L leads to a complexity growing (almost) polynomially as $L^{b/\log \log L}$ with some constant b . For higher W , as the tail decays more quickly, the dominating time shifts to an entanglement that appears more typically in the sample. In this regime, the computational time appears to scale linearly with L . In the case of a purely exponentially decaying tail in $P(S_E)$, the computation time experiences a Hagedorn-like transition at some W where the dominating computational cost switches from being in the tail of the distribution to some finite value near the peak of the distribution.

Thus far we have assumed a flat ES in order to make the approximation $\chi = e^{S_E}$. In reality, our ES decays on average quite quickly according to Eq 8, and the relationship between χ and S_E depends on the exact structure of the ES. However, χ should still scale exponentially with S_E on average. While the exact values of the coefficients

of the exponential will depend on the decay of the ES and the error threshold, the maximum bond dimension χ^{\max} will still scale as predicted. We confirm the dependence of χ on S_E by sampling the entanglement spectra obtained from the correlation matrix⁴⁷. The largest values in the many-body entanglement spectra are obtained from the eigenvalues of the single particle entanglement Hamiltonian \mathcal{H}_E via an iterative algorithm. We find the largest χ eigenvalues such that the total sum of the discarded spectrum is less than 10^{-5} . Figure 2(right) shows a histogram of χ vs. S_E at $W = 1.0$, clearly demonstrating a rough exponential relation between the two as expected. Notice that very large χ , much larger than reachable, is needed to attain this level of accuracy for low W . Figure 2(center) shows the max and mean χ as a function of L , demonstrating the expected scaling behaviour. At high W , χ becomes very small and the relationship between S_E and χ is more noisy.

V. FINDABILITY: HYBRID DMRG-X ALGORITHM

Having seen that the localized eigenstates of the XX model can be represented reasonably efficiently by MPSs, we turn to the issue of finding such representations using the DMRG-X method. In performing DMRG-X for this relatively simple model, we found that the algorithm was having great difficulty converging for certain disorder realizations within a particular range of parameters. The problem is not in the ability to represent such a state as an MPS, as better MPS representations can be constructed from the exact eigenstate, using an algorithm by S.R. White⁴⁹. White's algorithm iteratively applies pairwise rotations to an initial product state MPS, truncating the bond dimension as necessary, to construct a variational approximation—not necessarily the best one—to the exact eigenstate. Given such an MPS as an initial state, performing a DMRG-X sweep will typically worsen the state in the problematic parameter range. The typical energy variance σ_E^2 of eigenstates is shown for $L = 32$ as a function of W and χ in Figure 3(left). There is a particular region $1 < W < 2$ in which the algorithm fails to find the correct eigenstate. A significant portion of samples within this region did not show a monotone decrease in variance with increasing χ , meaning that the optimal state was never being reached.

A problem occurs when there is a near degeneracy between two many-body eigenstates that are not too many spin flips away. If the dimension of our restricted subspace is not sufficient to accurately capture the Hamiltonian acting on these two states, they may appear to be much closer in energy than they actually are, resulting in a false resonance in \mathcal{H}_{eff} . This happens as the subspace is optimized to accurately represent only the current state, not necessarily any others and therefore not the full Hamiltonian. The result of this is that some of the eigenstates of \mathcal{H}_{eff} are superpositions of the true

eigenstate with other nearby energy states. Thus, picking one of them based on maximal overlap does not move the overall state towards the correct answer. This is a problem in the findability.

Here, we present a small modification that can correct for this type of error. Rather than simply picking the eigenstate with highest overlap, we look at the n largest overlap eigenstates. The variance can then be optimized within this n -dimensional subspace, via a gradient method starting with the current state as a seed. Thus, variance and overlap are both taken into account in choosing the next state, which is closer to the true eigenstate than any eigenstate of \mathcal{H}_{eff} .

The variance calculation can be done efficiently in MPS language by keeping track of $\mathcal{H}_{\text{eff}}^2$ in addition to just \mathcal{H}_{eff} . Note that $\mathcal{H}_{\text{eff}}^2$ here is \mathcal{H}^2 in the effective subspace, not simply \mathcal{H}_{eff} squared. Then, the variance $\sigma_E^2 = |\langle \mathcal{H}_{\text{eff}}^2 \rangle - \langle \mathcal{H}_{\text{eff}} \rangle^2|$ can be optimized using a gradient based optimization algorithm such as the BFGS algorithm. These matrices can be further reduced to n by n matrices by only looking in the subspace of the n states with highest overlap with the previous state. An important implementation detail is that we optimize $\log \sigma_E^2$ rather than simply σ_E^2 , as we are concerned with very small values of the variance.

This hybrid algorithm works best when the MPS is already close to an eigenstate, as the overlap then provides a useful measure of which states to optimize over. It is therefore often helpful to run a few sweeps with $n = 1$ (the original algorithm) before increasing n . Rather than keeping n fixed, it is also possible to vary n . If we let $|\phi_i\rangle$ be the eigenstates of the effective Hamiltonian ordered from highest to lowest overlap with the previous state $|\psi\rangle$, then we can choose n to be the smallest integer satisfying

$$1 - \sum_{i=1}^n |\langle \phi_i | \psi \rangle|^2 < \epsilon \quad (10)$$

where ϵ is chosen to be a small number ($\epsilon = 10^{-10}$ is used in our calculations).

Figure 3(center) shows that the application of this hybrid algorithm greatly improves the variance near the troublesome region. The peak in the variance of the original algorithm arose due to convergence issues, which have been clearly resolved with the hybrid algorithm. Also shown are the variances obtained from the MPS constructed using White's algorithm^{49, 50}. Those variances do what one might expect from a direct truncation in the entanglement spectrum, in that they perform poorly at low W where there is high entanglement and monotonically do worse.

Indeed, the hybrid algorithm now always converges to a state, with variance decreasing regularly with χ . Exploring further, Figure 3 shows the mean entanglement entropy of the DMRG-X states in comparison with that of the exact eigenstates. We see that the state to which




FIG. 3. (left) Typical energy variance as a function of W for a system of size $L = 32$ for the original DMRG-X algorithm. Error bars show the standard error of the mean. DMRG-X works well at high W where the eigenstates are highly localized, and does worse as W is decreased. There is a region near $W = 1.5$ where the algorithm has the most difficulty converging onto an eigenstate. (center) Typical energy variance as a function of W for a system of size $L = 32$. The hybrid variance minimization algorithm is used, greatly improving convergence within the difficult region. The dashed lines show the variances of the compressed state obtained from the free fermion picture⁴⁹, which DMRG-X always outperforms. (right) The mean entanglement entropy obtained from DMRG-X using the hybrid algorithm for the $L = 32$ system. At low W , the obtained entanglement entropy is significantly lower than the exact values of the eigenstates.

DMRG-X converges towards is significantly lower in entanglement than the true exact eigenstates. This brings us to our next question: given that the hybrid DMRG-X algorithm does not suffer from problems of convergence (findability), how accurate are the states to the exact eigenstates?

VI. ACCURACY OF DMRG-X

We now turn to our goal, which is to evaluate the performance of DMRG-X on system sizes larger than accessible by many-body ED. To this end we use a modest bond dimension of $\chi = 32$, corresponding to a $2^2\chi^2 = 2^{12}$ dimensional effective Hamiltonian in the worst case. Thus the computational effort at the diagonalization step of DMRG-X is equivalent to the exact diagonalization of the full Hamiltonian of an $L = 12$ system. With this restriction we examine the accuracy of the algorithm on system sizes much bigger than $L = 12$ via the energy variance and, more revealingly, l-bit accuracy.

We begin by studying $L = 32$. We have seen in Fig 3(center) earlier that for $L = 32$, even in the low W regime where the localization length is larger than L (and significantly larger than the many-body ED length for $\chi = 32$), the algorithm is still able to converge to states to variances of $\sigma_E^2 \approx 10^{-6}$. Given the amount of entanglement in the low W region, one would have naively expected DMRG-X to perform very poorly at low W . The “flattening out” of the energy variance at low W and the lowered entanglement entropy is therefore DMRG-X performing better than expected — it is able to find a state significantly lower in entanglement while also maintaining a low energy variance. In this case, one may worry about how close the state is to the actual

eigenstates.

To better understand the quality of the state, Fig 4(left) shows the typical deviation of $\langle \tau_\alpha^z \rangle$ from ± 1 , defined by $\delta = 1 - |\langle \tau_\alpha^z \rangle|$. At $\chi = 32$, the typical l-bit expectation value is accurate up to roughly 10^{-6} at its worse. Figure 4(right) shows the average number of badly captured l-bits in a sample, which we define to be those with $\delta > 0.1$. Indeed, with high enough bond dimension DMRG-X is able to obtain most eigenstates with no bad l-bits, with the mean number of bad l-bits falling below 1. This is our central demonstration in this paper—that DMRG-X can obtain all of the l-bits correctly for system sizes bigger than those treatable by the ED routine utilized by the code. Evidently, if we were to use the actual maximum size treatable by ED ($L > 12$) we would expect to get to even bigger actual system sizes ($L > 32$) and still expect to get all of the l-bits correctly by our metric.

We now turn to the scaling with L while keeping $\chi = 32$ fixed. If the MPS exhibited a constant error density throughout, one would expect the energy variance to scale linearly with L . Figure 5(left) instead shows the variance scaling with L , which appears to be faster than linear. The typical l-bit errors, shown in Figure 5(right), are also very small and slowly increasing with L . The reason behind the super-linear growth in variance is that our states are in fact *not* of constant error density, as we are simply truncating each bond at χ singular values. To construct a state of constant error density, one would have to consider a variable χ on each bond such that the total discarded Schmidt weight is below an error threshold. In such a scheme one would still need to define a maximum allowable χ , χ_{\max} , such that the computation completes within a reasonable timeframe. One can then either discard samples that would exceed this




FIG. 4. (left) Typical value of δ for l-bits in a system of size $L = 32$, averaged over roughly 1000 disorder samples, with one eigenstate per sample. (right) Average number of l-bits with $\delta > 0.1$ in a system of size $L = 32$.

limit (leading to a low-entanglement bias) or truncate the bonds at χ_{\max} (which no longer produces constant error density states). Our fixed- χ choice corresponds to the latter with a zero error threshold.

Finally, a comment about the relationship between the two diagnostics we have used to gauge accuracy: the variance and the l-bit expectation values. These two measures are actually somewhat independent measures of accuracy. Expanding the variance in terms of l-bits,

$$\sigma_E^2 = \sum_{\alpha} h_{\alpha}^2 (1 - \langle \tau_{\alpha}^z \rangle^2) \quad (11)$$

$$+ \sum_{\alpha \neq \beta} h_{\alpha} h_{\beta} (\langle \tau_{\alpha}^z \tau_{\beta}^z \rangle - \langle \tau_{\alpha}^z \rangle \langle \tau_{\beta}^z \rangle) \quad (12)$$

one sees that the variance contains a sum over not only the l-bit expectation values, but also their correlators. Examining a typical DMRG state, there is a large degree of cancellation between the diagonal and off-diagonal contributions to the variance — thus leading the variance to be significantly lower than one would expect from only the l-bit expectation values. This is fairly easy to understand when considering the case where two l-bits α, β are very close in energy, $h_{\alpha} \approx h_{\beta}$. A state with $\langle \tau_{\alpha}^z \rangle = -\langle \tau_{\beta}^z \rangle = 1$ will have almost no energy difference from the state with the two of them flipped. One might worry that DMRG, being biased towards small entanglement, may be optimized by a superposition of these two states. In fact, going back to Figure 3(right), we see that our hybrid DMRG-X algorithm is remarkably able to obtain significant entanglement reduction in the low- W region, while still maintaining a low variance as well as very accurate l-bit expectation values (as can be seen in Fig 3(center) and Fig 4).

VII. CONCLUDING REMARKS

We have tested the performance of the DMRG-X algorithm for finding highly excited eigenstates of the disordered XX model. With some small improvements to the algorithm, we were able to go significantly beyond the range of system sizes accessible to exact diagonalization in the many-body Hilbert space. Most of the eigenstates are able to be obtained to high accuracy with only $\chi = 32$, using the l-bit expectation value as a metric. Thus, we have successfully shown that DMRG-X is able to go well beyond exact diagonalization in the localized phase, converging to eigenstates quickly with increasing χ .

There are still many ways of improving the capabilities of the algorithm. Rather than exact diagonalization of the effective Hamiltonian, shift-invert Lanczos can be used to obtain all the eigenstates within an energy window. The states with high overlap with the original state typically have very similar energy, so Lanczos is also able to find a relevant subspace of n states over which to optimize variance. This theoretically allows a great increase in the range of accessible χ .

With increased confidence in the algorithm in hand we can now return to studying the interacting problem. In addition to the transition to the ergodic phase in the model with interactions added to the XX chain, another potential application is to the phase transition between two localized phases, for example the spin glass to paramagnetic phase transition in the random transverse field Ising model.

ACKNOWLEDGMENTS

This research was supported by the Harvard Society of Fellows (VK), the DFG (Deutsche Forschungsgemeinschaft) Research Unit FOR 1807 through grants num-




FIG. 5. (left) Typical variance scaling with system size L at fixed bond dimension $\chi = 32$. The inset shows the same plot for $W = 5$ with fixed error density, truncating the Schmidt spectrum at a variable bond dimension such that the total cut-off error is less than 10^{-10} , which shows a linear scaling as expected. (right) Typical δ of l-bits scaling with system size L at fixed bond dimension $\chi = 32$. For high W , the typical l-bit is captured to machine precision, flattening out at $\langle \log_{10} \delta \rangle \approx 10^{-15}$

bers PO 1370/2-1 (FP), the NSF-DMR via grant number

1311781 (SLS) and the Alexander von Humboldt Foundation via a Humboldt Award (SLS).

-
- ¹ S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
 - ² U. Schollwock, *Annals of Physics* **326**, 96 (2011).
 - ³ I. P. McCulloch, arXiv:0804.2509 (2008).
 - ⁴ N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 030504 (2008).
 - ⁵ M. B. Hastings, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P08024 (2007).
 - ⁶ Z. L. I. Arad, A. Kitaev and U. Vazirani, arXiv:1301.1162 (2013).
 - ⁷ Although with possible exceptions, as discussed in Ref 51.
 - ⁸ Landau Zeph, Vazirani Umesh, and Vidick Thomas, *Nat Phys* **11**, 566 (2015).
 - ⁹ P. W. Anderson, *Phys. Rev.* **109**, 1492 (1958).
 - ¹⁰ D. M. Basko, I. L. Aleiner, and B. L. Altshuler, *Annals of Physics* **321**, 1126 (2006).
 - ¹¹ A. Pal and D. A. Huse, *Phys. Rev. B* **82**, 174411 (2010).
 - ¹² V. Oganesyan and D. A. Huse, *Phys. Rev. B* **75**, 155111 (2007).
 - ¹³ R. Nandkishore and D. A. Huse, *Annual Review of Condensed Matter Physics* **6**, 15 (2015).
 - ¹⁴ E. Altman and R. Vosk, *Annual Review of Condensed Matter Physics* **6**, 383 (2015).
 - ¹⁵ B. Bauer and C. Nayak, *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P09005 (2013).
 - ¹⁶ D. Pekker and B. K. Clark, *Phys. Rev. B* **95**, 035116 (2017).
 - ¹⁷ A. Chandran, J. Carrasquilla, I. H. Kim, D. A. Abanin, and G. Vidal, *Phys. Rev. B* **92**, 024201 (2015).
 - ¹⁸ D. A. Huse, R. Nandkishore, V. Oganesyan, A. Pal, and S. L. Sondhi, *Phys. Rev. B* **88**, 014206 (2013).
 - ¹⁹ D. Pekker, G. Refael, E. Altman, E. Demler, and V. Oganesyan, *Phys. Rev. X* **4**, 011052 (2014).
 - ²⁰ S. A. Parameswaran, A. C. Potter, and R. Vasseur, ArXiv e-prints (2016), arXiv:1610.03078 [cond-mat.dis-nn].
 - ²¹ V. Oganesyan and D. A. Huse, *Phys. Rev. B* **75**, 155111 (2007).
 - ²² J. A. Kjäll, J. H. Bardarson, and F. Pollmann, *Phys. Rev. Lett.* **113**, 107204 (2014).
 - ²³ R. Vosk, D. A. Huse, and E. Altman, *Phys. Rev. X* **5**, 031032 (2015).
 - ²⁴ A. C. Potter, R. Vasseur, and S. A. Parameswaran, *Phys. Rev. X* **5**, 031033 (2015).
 - ²⁵ T. Devakul and R. R. P. Singh, *Phys. Rev. Lett.* **115**, 187201 (2015).
 - ²⁶ D. J. Luitz, N. Laflorencie, and F. Alet, *Phys. Rev. B* **91**, 081103 (2015).
 - ²⁷ L. Zhang, B. Zhao, T. Devakul, and D. A. Huse, *Phys. Rev. B* **93**, 224201 (2016).
 - ²⁸ M. Serbyn and J. E. Moore, *Phys. Rev. B* **93**, 041424 (2016).
 - ²⁹ L. Zhang, V. Khemani, and D. A. Huse, *Phys. Rev. B* **94**, 224202 (2016).
 - ³⁰ V. Khemani, D. N. Sheng and D. A. Huse, arXiv:1702.03932 (2017).
 - ³¹ T. Devakul and R. R. P. Singh, *Phys. Rev. Lett.* **115**, 187201 (2015).
 - ³² D. J. Luitz, N. Laflorencie, and F. Alet, *Phys. Rev. B* **91**, 081103 (2015).
 - ³³ V. Khemani, S. P. Lim, D. N. Sheng, and D. A. Huse, ArXiv e-prints (2016), arXiv:1607.05756 [cond-mat.dis-nn].
 - ³⁴ T. Grover, ArXiv e-prints (2014), arXiv:1405.1471 [cond-mat.dis-nn].
 - ³⁵ F. Pollmann, V. Khemani, J. I. Cirac, and S. L. Sondhi, *Phys. Rev. B* **94**, 041116 (2016).
 - ³⁶ T. B. Wahl, A. Pal, and S. H. Simon, ArXiv e-prints

- (2016), [arXiv:1609.01552 \[cond-mat.dis-nn\]](#).
- ³⁷ V. Khemani, F. Pollmann, and S. L. Sondhi, *Phys. Rev. Lett.* **116**, 247204 (2016).
- ³⁸ X. Yu, D. Pekker, and B. K. Clark, *Phys. Rev. Lett.* **118**, 017201 (2017).
- ³⁹ S. P. Lim and D. N. Sheng, *Phys. Rev. B* **94**, 045111 (2016).
- ⁴⁰ M. Serbyn, A. A. Michailidis, D. A. Abanin, and Z. Papić, *Phys. Rev. Lett.* **117**, 160601 (2016).
- ⁴¹ D. A. Huse, R. Nandkishore, and V. Oganesyan, *Phys. Rev. B* **90**, 174202 (2014).
- ⁴² M. Serbyn, Z. Papić, and D. A. Abanin, *Phys. Rev. Lett.* **111**, 127201 (2013).
- ⁴³ G. Vidal, *Phys. Rev. Lett.* **91**, 147902 (2003).
- ⁴⁴ Kappus M. and Wegner F., *Zeitschrift für Physik B Condensed Matter* **45**, 15 (1981).
- ⁴⁵ They are obviously local if one takes the definition of an MBL system to be the existence of a finite depth unitary transformation that diagonalizes the Hamiltonian.
- ⁴⁶ The argument in Ref⁴⁰ relies on the l-bit operators flipping many spins within a radius r , but in a noninteracting model the l-bit operators only flip two.
- ⁴⁷ I. Eisler, *Journal of Physics A: Mathematical and Theoretical* **42**, 504003 (2009).
- ⁴⁸ Finding the largest k eigenvalues of ρ_A involves finding the k lowest total eigenenergies given all the single particle energies, which is a standard computational problem.
- ⁴⁹ M. T. Fishman and S. R. White, *Phys. Rev. B* **92**, 075132 (2015).
- ⁵⁰ Note that it is not that the algorithm of Ref⁴⁹ is failing at low W . Similar curves were obtained for smaller L accessible by ED, and a standard MPS compression of the ED state yielded similar results.
- ⁵¹ J. Eisert, *Phys. Rev. Lett.* **97**, 260501 (2006).
- ⁵² S. Aubry and G. Andre, *Ann. Isr. Phys. Soc.* **3**, 133 (1980).

A. LOCALITY OF L-BIT OPERATORS

We have a Jordan Wigner transformed spin chain, diagonalized by a set of fermionic operators

$$\mathcal{H} = \sum_{\alpha} \epsilon_{\alpha} a_{\alpha}^{\dagger} a_{\alpha} \quad (13)$$

where each $a_{\alpha}^{\dagger} = \sum_i u_{\alpha}(i) c_i^{\dagger}$, and in the spin language, $c_i^{\dagger} = \prod_{j < i} (-1)^{(S_j^z + \frac{1}{2})} S_i^{+}$, and similarly for the conjugates. We want to show that the bosonic raising and lowering operators τ_{α}^{+} and τ_{α}^{-} are also localized, given that the fermionic ones a_{α}^{\dagger} and a_{α} are in terms of the physical c_i^{\dagger}, c_i operators.

The corresponding bosonic raising operator for $\tau_{\alpha}^z = 2a_{\alpha}^{\dagger} a_{\alpha} - 1$ is then defined by an inverse Jordan Wigner transformation

$$\begin{aligned} \tau_{\alpha}^{+} &= \prod_{\beta=1}^{\alpha-1} (-1)^{a_{\beta}^{\dagger} a_{\beta}} a_{\alpha}^{\dagger} \\ &= (-1)^{\sum_{\beta=1}^{\alpha-1} a_{\beta}^{\dagger} a_{\beta}} \sum_{i=1}^L u_{\alpha}(i) (-1)^{\sum_{j=1}^{i-1} c_j^{\dagger} c_j} S_i^{+} \end{aligned} \quad (14)$$

For simplicity, we will assume that $u_{\alpha}(i)$ is strictly localized within a range ξ , and is therefore only nonzero if

$|\alpha - i| \leq \xi$ (in reality, there will be an exponential decay with lengthscale ξ , and our argument can be generalized for some error threshold). We can then write τ_{α}^{+} as

$$\begin{aligned} \tau_{\alpha}^{+} &= (-1)^{\sum_{\beta=1}^{\alpha-1} a_{\beta}^{\dagger} a_{\beta}} (-1)^{\sum_{j=1}^{\alpha-\xi-1} c_j^{\dagger} c_j} \\ &\quad \sum_{i=\alpha-\xi}^{\alpha+\xi-1} u_{\alpha}(i) (-1)^{\sum_{j=\alpha-\xi}^{i-1} c_j^{\dagger} c_j} S_i^{+} \end{aligned} \quad (16)$$

Expressing the first exponent in Equation 16 as

$$\sum_{\beta=1}^{\alpha-1} a_{\beta}^{\dagger} a_{\beta} = \sum_{j,k=1}^L \sum_{\beta=1}^{\alpha-1} u_{\beta}(j) u_{\beta}^{*}(k) c_j^{\dagger} c_k \quad (17)$$

we see that if $j < \alpha - \xi$ or $k < \alpha - \xi$, then the sum over β contains all non-zero elements of $u_{\beta}(i) u_{\beta}^{*}(j)$, resulting in a dirac delta $\delta_{i,j}$ by completeness. This reduces the expression to

$$\begin{aligned} \sum_{\beta=1}^{\alpha-1} a_{\beta}^{\dagger} a_{\beta} &= \sum_{j=1}^{\alpha-\xi-1} c_j^{\dagger} c_j + \\ &\quad \sum_{j,k=\alpha-\xi}^{\alpha+\xi-1} \sum_{\beta=1}^{\alpha-1} u_{\beta}(j) u_{\beta}^{*}(k) c_j^{\dagger} c_k \end{aligned} \quad (18)$$

where we have restricted the summation indices $j, k \leq \alpha + \xi - 1$, since $\beta \leq \alpha - 1$ and we have assumed $u_{\beta}(i)$ is zero for $i > \beta + \xi$.

The two sums in Eq 18 commute, and therefore the $c_j^{\dagger} c_j$ sum can be canceled out with the chain in Eq 16, leaving

$$\tau_{\alpha}^{+} = (-1)^{\sum_{j,k=\alpha-\xi}^{\alpha+\xi-1} \sum_{\beta=1}^{\alpha-1} u_{\beta}(j) u_{\beta}^{*}(k) c_j^{\dagger} c_k} \quad (19)$$

$$\sum_{i=\alpha-\xi}^{\alpha+\xi} u_{\alpha}(i) (-1)^{\sum_{j=\alpha-\xi}^{i-1} c_j^{\dagger} c_j} S_i^{+} \quad (20)$$

which is strictly localized within $\alpha \pm \xi$.

B. MPO REPRESENTATION OF L-BIT OPERATORS

The $a_{\alpha}^{\dagger} a_{\alpha}$ operators can be constructed as an MPO by hand. The methodology is similar to that used to construct the MPO representation of the Hamiltonian². We use the internal dimensions $\{s_i\}$ to keep track of the different terms in the expansion of $a_{\alpha}^{\dagger} a_{\alpha}$. We can define the full operator as an MPO,

$$\begin{aligned} &\langle \{s'_i\} | a_{\alpha}^{\dagger} a_{\alpha} | \{s_i\} \rangle \\ &= \sum_{\{s_i\}} \delta_{s_i,1} \delta_{s_{L+1},4} \prod_{i=1}^L \langle s'_i | M_{s_i, s_{i+1}}^{[i]} | s_i \rangle \end{aligned} \quad (21)$$

where each s_i goes from 1 to 4 and $\delta_{i,j}$ is the Dirac delta.

The different terms in $a_\alpha^\dagger a_\alpha = \sum_{i,j} u_\alpha(i) u_\alpha^*(j) c_i^\dagger c_j$ can be taken into account by defining $M_{s_i, s_{i+1}}^{[i]}$ appropriately:

$$M^{[i]} = \begin{bmatrix} \mathbb{1} & u_\alpha(i)\sigma^+ & u_\alpha^*(i)\sigma^- & |u_\alpha(i)|^2(\sigma^z + 1)/2 \\ 0 & -\sigma^z & 0 & u_\alpha^*(i)\sigma^- \\ 0 & 0 & -\sigma^z & u_\alpha(i)\sigma^+ \\ 0 & 0 & 0 & \mathbb{1} \end{bmatrix} \quad (22)$$

Thus, the operator $a_\alpha^\dagger a_\alpha$ has been expressed as an MPO of internal dimension 4, and the expectation value in an MPS can be efficiently calculated. One can then simply take $\tau_\alpha^z = 2a_\alpha^\dagger a_\alpha - 1$.

C. GRIFFITHS EFFECTS AND QUASIPERIODICITY

Many of the difficulties in studying disordered systems using DMRG arise due to the existence of locally clean-looking patches, which are what the tail of the distribution $P(S_E)$ consist of. An exponentially decaying tail in $P(S_E)$ will lead to the typical maximum entanglement in a system of size L growing as $S_E^{\max} \sim \log L$, resulting in the computational difficulty of DMRG grow faster than linearly. To avoid the difficulties of these Griffiths region in studies of localization, one possibility is to examine quasiperiodic systems.

One commonly studied possibility is the Aubry-André model, which is given by the Hamiltonian

$$\mathcal{H}_{AA} = - \sum_i \left(c_i^\dagger c_{i+1} + \text{h.c.} \right) + W \sum_i V(i) c_i^\dagger c_i \quad (23)$$

where $V(i) = \cos(2\pi\omega(i + \phi))$, ω is an irrational number which we choose to be 1 over the golden ratio, and ϕ is some offset. The exact delocalization-localization critical point of this model is known to be at $W = 2$, and

the critical wavefunctions can be found exactly⁵². The quasiperiodic potential does not allow for large Griffiths-like regions to exist.

Despite this, sampling chains of length L (over different phase offsets), there is still a clear logarithmic growth of S_E^{\max} with L , corresponding to an exponentially decaying tail in $P(S_E)$, which is reminiscent of the existence of Griffiths regions!

This effect can be understood as coming from the existence of special reflection symmetry points in the quasiperiodic potential. For example, if $\phi = 0$, all the sites on the left side of $i = 0$ will be in exact resonance with another site on the right side, $V(-i) = V(i)$. Each resonant pair will then form an odd-even superposition and the many-body entanglement entropy across that cut will grow approximately as $S_E^{\max} \sim L$. Thus, there exist special values of ϕ for which the entanglement diverges across a cut.

In an actual sample, this exact resonances will not occur, only approximate resonances. Each sample can be thought of as sampling ϕ L times, and so one will typically be at most $\sim 1/L$ away from the exact resonance point, and sites will be detuned by $\sim 1/L$. The effective hopping between two sites distance x apart scales as $\sim e^{-x/\xi}$ for localization length ξ . Therefore, for such a near resonance, only pairs of sites within range x such that $e^{-x/\xi} \gtrsim 1/L$ will form superpositions. This leads to an entanglement going as $S_E^{\max} \sim x \sim \log L$. Thus, a completely different mechanism in the quasiperiodic model is allowing for logarithmic growth in maximum entanglement.

To avoid a logarithmic growth of S_E^{\max} , one should then avoid a quasiperiodic potential with reflection symmetry. This is possible by a potential such as $V(i) = \cos(2\pi\omega(i + \phi)) + \sin(4\pi\omega(i + \phi))$, which does not exhibit any even reflection symmetry. The S_E distribution for this model has a very sharp cutoff and the maximum entanglement therefore does not keep growing as $\log L$.